

performing a direct look-up of a first phoneme of the word in a database of classified words organized alphabetically by their phonemes,

matching the first phoneme of the word to a first word in the database having the same first phoneme,

selecting the first word in the database having the same first phoneme and following words in the database to form a sub-list,

matching the first phoneme of the word to the first phoneme of the classified words in the sub-list to form a class list, and

placing the word in the class list.

REMARKS

Claims 1-14 are pending with claims 1, 7, 8, and 10 being independent. Applicants acknowledge the Examiner's allowance of claims 1-6. The Examiner has rejected claim 7 as anticipated by Brown, and has rejected claims 8-14 as anticipated by Hutchins. We first address the rejection over Hutchins.

Claims 8 and 10

Claims 8 and 10, as amended, are directed to a method of assigning a pre-filtering class to a word when adding a word to a speech recognition dictionary. As described at page 36, lines 3-7, once the phonetic spelling of a new word has been generated, the new word is added to the speech recognition dictionary. To determine where in the dictionary to locate the

new word, the new word is assigned a pre-filter class by comparing the new word to words already classified.

The method of claim 8 includes assigning a pre-filtering class to a word by matching first two letters of the word to classified words starting with the same two letters to form a sub-list of classified words, matching phonemes of the word to phonemes of the classified words in the sub-list to form a class list, and placing the word in the class list.

The method of claim 10 includes assigning a pre-filtering class to a word by performing a direct look-up of a first phoneme of the word in a database of classified words organized alphabetically by their phonemes, matching the first phoneme of the word to a first word in the database having the same first phoneme, selecting the first word in the database having the same first phoneme and following words in the database to form a sub-list, matching the first phoneme of the word to the first phoneme of the classified words in the sub-list to form a class list, and placing the word in the class list.

Applicants submit that Hutchins does not describe assigning a pre-filtering class to a word when adding a word to a speech recognition dictionary, as claimed.

Hutchins is primarily concerned with a method for speech recognition. The method is based on SubSyllable "spellings." Hutchins defines a SubSyllable as "a cohesive segment of the acoustic signal that is generally smaller than a phoneme" (col. 5, lines 23-25).

In Hutchins, to recognize incoming speech, the speech is digitized and the digitized speech is broken into SubSyllables (see, e.g., col. 13, lines 46-67). The "spelling" of a SubSyllable is a phonetic spelling based on symbology conventional in linguistics (see, e.g., col. 14, lines 5-22).

Hutchins achieves recognition from a string of SubSyllables by comparing the string of SubSyllables "to a SubSyllable Grammar 24, which contains tables of predetermined 'spellings' of words or parts of words.... If a sequence of SubSyllable spellings from successive frames matches a spelling in the table, recognition is achieved" (col. 14, lines 43-52).

Hutchins describes adding words to a vocabulary by writing SubSyllable spellings to describe them (see col. 24, line 50 to col. 27, line 16). However, Hutchins does not describe adding a word to a speech recognition dictionary by assigning a prefiltering class to a word.

In rejecting claims 8 and 10, the Examiner has referred to various sections of columns 10, 11, and 15. These sections are directed to Hutchins' method of speech recognition and are not directed to assigning a pre-filtering class to a word when adding a word to a speech recognition dictionary.

Furthermore, Hutchins' method of adding a word to a vocabulary does not address the placement of the word in a speech recognition dictionary. There is no discussion in Hutchins of assigning a pre-filtering class to a word when adding a word to a dictionary, as recited in claims 8 and 10. Nor does Hutchins describe matching aspects of the word to aspects of words that

have already been classified. In particular, Hutchins does not describe matching the first two letters of a word with classified words and matching phonemes of the word with classified words, as recited in claim 8, nor does Hutchins describe matching the first phoneme of the word to words in a database having the same first phoneme, as recited in claim 10.

Therefore, applicants submit that claims 8 and 10, and the claims dependent thereon, are patentable over Hutchins for at least these reasons.

Claim 7

We now turn to the Examiner's rejection of claim 7 as anticipated by Brown.

Claim 7 is directed to adding a word to a speech recognition vocabulary. A spelling of the word is used to generate a net of possible phonetic pronunciations of the word. The possible pronunciation from the net that best matches an utterance of the word is selected using speech recognition. An example of a net generated according to the invention is shown in Fig. 16 and described on page 35. By generating a net, all phonemes associated with a letter node of the net can be evaluated with paths of the net that do not contain phonemes corresponding to the spoken word being pruned out using speech recognition.

Applicants submit that Brown does not describe using the spelling of a word to generate a net of possible pronunciations of a word.

Brown describes using the spelling of a word to create a list of possible word models using a model of spelling to sound probabilities, such as the model shown in TABLE 2 of Brown. In order to limit the number of possible word models in the list, Brown selects only the one or two sound models having the highest probabilities (see, e.g., col 6, lines 7-11). The best matched possible word model is then determined by a weighted combination of (a) the closeness of a match between the probabilistic word model and the values of the acoustic feature of each utterance, and (b) the closeness of a match between the probabilistic word model and the spelling of the word (see, e.g., col. 6, lines 48-54).

However, Brown does not describe using the spelling of a word to generate a net of possible phonetic pronunciations of a word, as claimed. Rather than using a net, which has the advantage of permitting the evaluation of all possible phonemes and of managing the large number of possible phoneme matches by using speech recognition to prune out paths of a net that do not contain phonemes corresponding to the spoken word, Brown limits the number of possible word models by selecting only the highest probability sound models for evaluation. Brown's method has the drawback of never analyzing the lower probability sound models. Applicants' method addresses this very drawback.

In rejecting claim 7, the Examiner has referred to various sections of columns 12 and 13 of Brown. These sections are directed to Brown's method of creating an expansion set of word models. How closely an expansion set of word models matches

an uttered word is compared to how closely a selection set of word models matches the uttered word to determine which of the sets is used to model the word. These sections of Brown are not directed to using the spelling of a word to generate a net.

Therefore, applicants submit that claim 7 is patentable over Brown for at least these reasons.

Applicants submit that all of the claims are now in condition for allowance, which action is requested.

Please charge any fees, or make any credits, to Deposit Account No. 06-1050.

Respectfully submitted,

Date:

June 4, 1999



Phyllis K. Kristal
Reg. No. 38,524

Fish & Richardson P.C.
225 Franklin Street
Boston, MA 02110-2804

Telephone: 617/542-5070
Facsimile: 617/542-8906

357160.B11